

Stochastic Simulation of Inherited Kinship-Driven Altruism

Heather Turner; Dimitar Kazakov

Computer Science Department

University of York

{hrt103,kazakov}@cs.york.ac.uk

1 Introduction

The aim of this research is to assess the rôle of a hypothetical inherited feature (gene) promoting altruism between relatives as a factor for survival. The two main goals are, firstly, to replicate the phenomenon of altruism, which has been observed in nature, and show that the proposed mechanism leads to altruistic individuals being selected by evolution. Secondly, the research aims to provide an implementation of a Multi-Agent System (MAS) employing a model of natural selection, which is different from the one commonly used in Computer Science (Goldberg, 1989), and, hopefully, closer to the one existing in nature.

Altruism can be defined as selfless behaviour, action that will provide benefit to another at no gain to the actor himself, and possibly even to his detriment. In *kinship-driven altruism*, this behaviour is directed between individuals who are related. Hamilton (1964a) introduces an analytical model, in which altruistic behaviour towards relatives is favoured by evolution, provided that the amount of help that an individual bestows on relatives of a given distance is appropriately measured.

Both MASs (Wooldridge and Jennings, 1995) and Genetic Algorithms (GAs) (Goldberg, 1989) can be used effectively to simulate the interaction of a population that evolves over a period of time. A MAS allows study of the interactions at the level of the individual, while a GA is a better tool for generalisation over an entire population. In a GA, no distinction is made between individuals with the same *genotype* (i.e., inherited features), whereas in a MAS these are represented by different *phenotypes*, or set of observable characteristics resulting from the interaction of each genotype with the environment (Thompson, 1996). The use of MAS with large populations is limited by the requirement for extra resources to represent individual phenotypes. In a GA, the individual is anonymous, so there is no capacity to “zoom-in” on its behaviour, but in contrast, there is the possibility of scaling up to consider a much larger population, which may be statistically more relevant.

The GA uses a fitness function to estimate how well

each individual will fare in the future and uses this to influence the likelihood that they survive to subsequent generations. A MAS uses information about the current position of an individual in the environment, and taking into account its internal state, considered to be the cumulative result of its actions and experiences in the past, determines its actions. In a GA, the population size is fixed, and during each system cycle, individuals may be selected to mate (and be replaced by their descendants) or they pass to the next generation. The anonymity of each individual is suited to the probabilistic selection afforded to this algorithm, and the resulting possibility that clones of an individual be produced in future generations. Without this anonymity, in a system that “tracks” the behaviour of individuals through the generations, complications could arise on cloning. Attachment of energy values becomes difficult if the probabilistic freedom is to be maintained without producing a system that can produce and destroy energy at will. In a MAS, the population size is not explicitly constrained, and the internal state of an individual determines its lifespan. A system cycle will not generally represent an entire generation, as individuals may survive for many cycles. Table 1 summarises the main differences between the GA and MAS models of natural selection.

We combine features of each approach to produce a more scalable, personality-driven system without a modelled spatial dimension. The probabilistic nature of all events and the high level of abstraction typical for the GA are preserved. However, the description of each individual consists of a set of inherited features (genome) along with a—very abstract—description of the actual organism (phenotype). The internal state of each individual is changed by the interaction with a very simple, abstract, environment, in which both the selection of an individual’s action and its outcome are modelled as probabilistic functions. This permits to avoid the use of an explicit fitness function, and describe the survival of an individual directly as a probabilistic function of its internal state (e.g., current energy levels).

Our system is designed to simulate a population in which some of the individuals are carriers of a gene forc-

Table 1: MAS vs GA simulation of natural selection

Feature	MAS	GA
Representation of individuals	genotype + phenotype	genotype only
Survival of individuals	deterministic, based on the lifetime interaction with environment	probabilistic, based on genotype's fitness
Population size	unlimited	fixed
Environment resources	limited capacity	use bounded by max. pop. size
Preservation of energy	enforced	not considered

ing them to share their energy with the individuals they meet in proportion to the degree of kinship (i.e., number of shared genes). The exact sharing policy is subjected to selection and studied. The altruistic gene is passed with a certain probability from parent to child. Food consumption results in increased individual energy level. Falling below a certain energy level means death. An encounter of two individuals of the same species could result in the creation of an offspring if their energy levels are sufficient. The initial energy level of the offspring is subtracted from that of the parents.

This research uses the hybrid platform described above to study from a different angle an extended version of some of the experiments with kinship-driven altruism performed by Barton (2001). As in Barton's study, natural selection works to produce an optimum sharing function when it is left to vary.

2 Altruism and Darwinian Theory

The possible evolution of a selfless gene is an interesting area of study as it does not necessarily seem intuitive that an individual should value the survival of another to the extent of causing detriment to itself (perhaps by decreasing his own chance of mating or survival) in order to help the other. It would be in contrast to the classic Darwinian theory of natural selection, according to which selfish individuals would always take the upper hand, and eliminate altruists, as the behaviour of the latter would by definition hinder their reproductive success. There is evidence however, as Hamilton (1964b) illustrates, that many species in nature exhibit this altruistic trait. Neo-Darwinian theory (Watson, 1995) attempts to provide an explanation with the idea of 'inclusive fitness', and the hypothesis that natural selection works not at the level of an individual, but on each individual gene. Many individuals can carry copies of the same gene, and if these individuals could identify one another, it would be pos-

sible for them to aid in the process of natural selection over that gene by attempting to secure reproductive success and the passing of this gene to the next generation. The evidence provided by Hamilton suggests that nature has evolved to recognise that it is likely for close relatives to have similar genetic makeup. In Hamilton's model, the degree of kinship is quantified, and it can then be used to determine how much help an individual can bestow on a relative, at detriment to itself and yet still be likely to benefit the inclusive fitness, the 'fitness' of the gene.

Barton (2001) used a MAS to model a population of individuals who behaved altruistically competing in an environment with a population of the same size that was not altruistic. His MAS used GA principles by associating genes with each individual in an attempt to find optimum solutions for variables used in his simulations. In some of his experiments, it was the sharing population that prevailed, in others, the non-sharing population over-ran the environment. He quotes 'Gause's Competitive Exclusion Principle', stating 'no two species can coexist if they occupy the same niche', and hypothesises that given the limitations of his simulated system, his competing populations are likely to 'end up having the same, or very similar, niches'.

In the MAS he uses, there are agents to represent food and the individuals of each population. The environment is represented on a grid with varying terrain that could restrict movement, or provide water as sustenance to fulfil 'thirst,' one of the 'drives' that describe the internal state of an agent in a given cycle. Each agent uses the values of its drives, its immediate surroundings and some deterministic rules to make life choices in each cycle.

3 Design

The system we have implemented to investigate altruistic behaviour combines features used in a MAS and those used in a GA. Rather than providing co-ordinates for the position of each individual in the system, we model encounters with food (energy) and other individuals probabilistically, reflecting the likelihood that these would occur in a given cycle. We do not constrain the population size, thus permitting easier comparisons with Barton's work (Barton, 2001). We stem the growth of our population by increasing the probability of random death as the individual ages. The individuals in our implementation retain a portion of genetic material, encoding their behaviour, and sharing policy, and thus allowing evolution of optimum policies. The diagram provides the proposed environmental interaction module for our system. Each individual stores as its phenotype the value of its sex drive, its hunger (or energy level), age and the probability of survival. These values are updated in each system cycle. Figure 1 contains an outline of the proposed simulation, where individual boxes have the following functions:

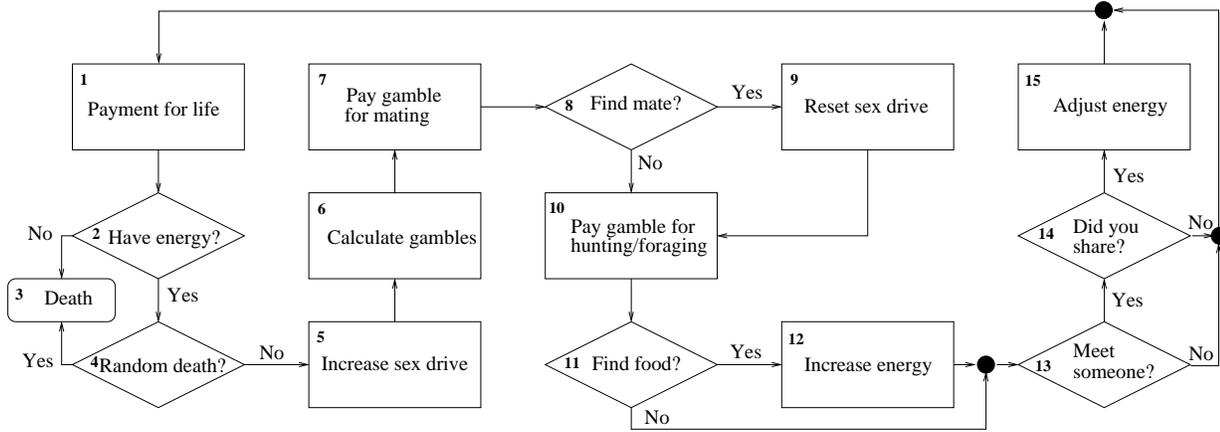


Figure 1: Simulation outline

1. Make a payment of energy to the environment (energy expended to survive generation).
2. If all energy is used up, one dies.
3. Individual has ‘died’ and is therefore removed from the population.
4. Random death occurs with some probability for each individual (probability increases with age).
5. Increase sex drive, and thus priority of reproduction.
6. Genetic material encodes a function to determine behaviour based on the values of the drives. This function produces “gambles” dictating how much of the available energy to expend in search for a mate or food, if any.
7. The gamble for mating is ‘paid’ to the system.
8. A probability distribution over the gambles is used to choose individuals for the mating pool(s). Pairs selected at random from the mating pool are deemed to have ‘met’ with some probability. Each must satisfy certain energy requirements. The probability that they mate is calculated from their sex drives and determines whether or not they actually ‘mate’. On mating, new individuals are created from clones of the genetic material, and by resetting non-genetic parameters. Each parent contributes energy for sharing equally amongst the offspring. The clones undergo crossover producing two children to be included in the population for the next cycle.
9. The sex drive of the individuals who mated successfully is reset.
10. The gamble for hunting/foraging is ‘paid’ to the system.
11. A probability distribution based on the gamble determines how much energy an individual receives. For a gamble of zero, the probability that an individual receive any energy should be very low.

12. Energy level is increased by the amount of food found.
13. Pairs are further selected from the population, and with some probability are deemed to meet.
14. If the better fed of the pair is an altruist, they decide to share as per his genetically encoded sharing policy.
15. The energy of each individual is then adjusted as appropriate.

4 Experiments and Evaluation

The gambling policies of the tool described in the previous section were implemented as stochastic functions of the current needs of the individual. A sigmoid function was used to provide a nonlinear mapping from gamble to average win μ . The actual win of each gamble was generated according to a Gaussian distribution $G(\mu, \sigma)$ where the ratio σ/μ was kept constant for all μ , to ensure that only a very small, fixed proportion of the wins were negative; these, when generated, were reset to zero.

In the simulation, spatial phenomena (food discovery, encounter with another individual) are represented as random processes with a certain probability. In this case, physical distance between individuals is ignored, and the encounter of each pair is equally probable.

The so specified tool was implemented in C++, and used to study the influence of several factors on the evolution of altruistic behaviour. In all cases, the evaluation assesses whether the hypothetical altruistic gene is selected by evolution, and study the circumstances in which this happens.

Degree of kinship Individuals may (1) have a complete knowledge of their genealogy (*Royalty* model), (2) estimate the degree of kinship according to the presence of some inherited visible indicators (*Prediction*), or (3) not have this information available (*Unknown*).

Type of sharing function Three social models are considered. *Communism* equalises the energy levels of two individuals with the same genome. *Progressive taxation with a non-taxable allowance* is a simple linear function with a threshold: $y = \alpha(x - \theta)$ for $x > \theta$; $y = 0$ otherwise. *Poll tax* defines an altruistic act between two individuals as an exchange of a fixed amount of energy pt set in the genes of the donor, which does not depend on the energy level of either individual. The above descriptions corresponds to the case of sharing between two individuals with the same set of genes. In all other cases, the actual amount given is reduced in proportion to the difference between the two individuals' genomes, as derived from the perceived degree of kinship.

All combinations of the above two factors have been studied by running three times each of the nine possible experiments (see Figure 2). For technical reasons, the Royalty model of kinship stored only relatives up to first cousins (inclusive). In the predictive mode, all genes but (the altruistic) one were made visible to other individuals, and the difference between two sets of genes was defined according to a simple linear metric.

All parameters of the sharing functions (α, θ) , resp. pt were initially set at random, and let to evolve. When employing the Unknown model of kinship, the most optimistic assumption was made, i.e., the donor treated the aid receiver as an identical twin brother.

In the experiments, all individuals carry a gene defining them as either selfish or altruistic. Simply counting the individuals carrying either gene is a good measure of the altruism in the population only in a communist society. In the other two cases, individuals, which are nominally altruistic, can have their sharing parameters set in a way, which reduces the effects of altruism to an arbitrary low level, e.g., α or $pt \rightarrow 0$, $\theta \rightarrow \infty$. In these cases, the ratio of what is given to what is actually owned by the individual, integrated over the whole energy (food) range, is considered a more appropriate measure. The idea in the case of progressive taxation is shown in Figure 2 where a nominally altruistic individual is assigned a degree of altruism given by the ratio of the filled triangle and the square made of the ranges of energy owned and exchanged.

The graphs in Figure 2 are self-explanatory. In brief, the use of either perfect knowledge of the degree of kinship or a sharing function based on progressive taxation ensures that a substantial level of altruism is selected and maintained in the population. In the remaining cases, altruists perish faster when less kinship knowledge is available. The population size remains the same in all cases, and is given by the amount of food supplied. A representative example of the way in which the population size evolved is shown in Figure 4 on the case of Royalty with Progressive Taxation.

Initial ratio between altruistic and selfish individuals
To study the influence that the initial proportion of altru-

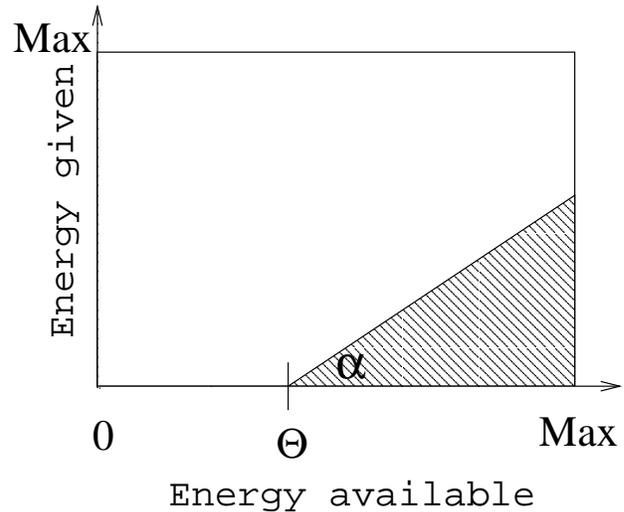


Figure 2: Measure of altruism

istic to selfish individuals has on the levels of altruism selected by evolution, the Royalty with Progressive Taxation experiment was run with several initial values for this ratio. The results in Figure 3 show that the system reaches a dynamic equilibrium which, in the cases shown, does not depend on the initial point.

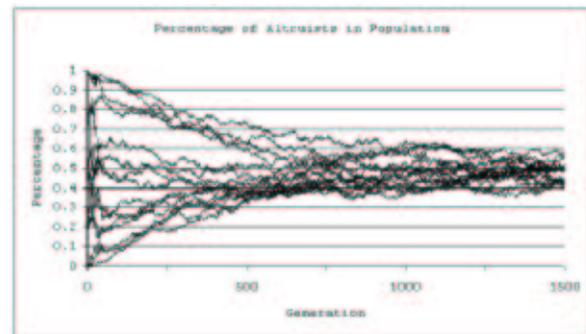
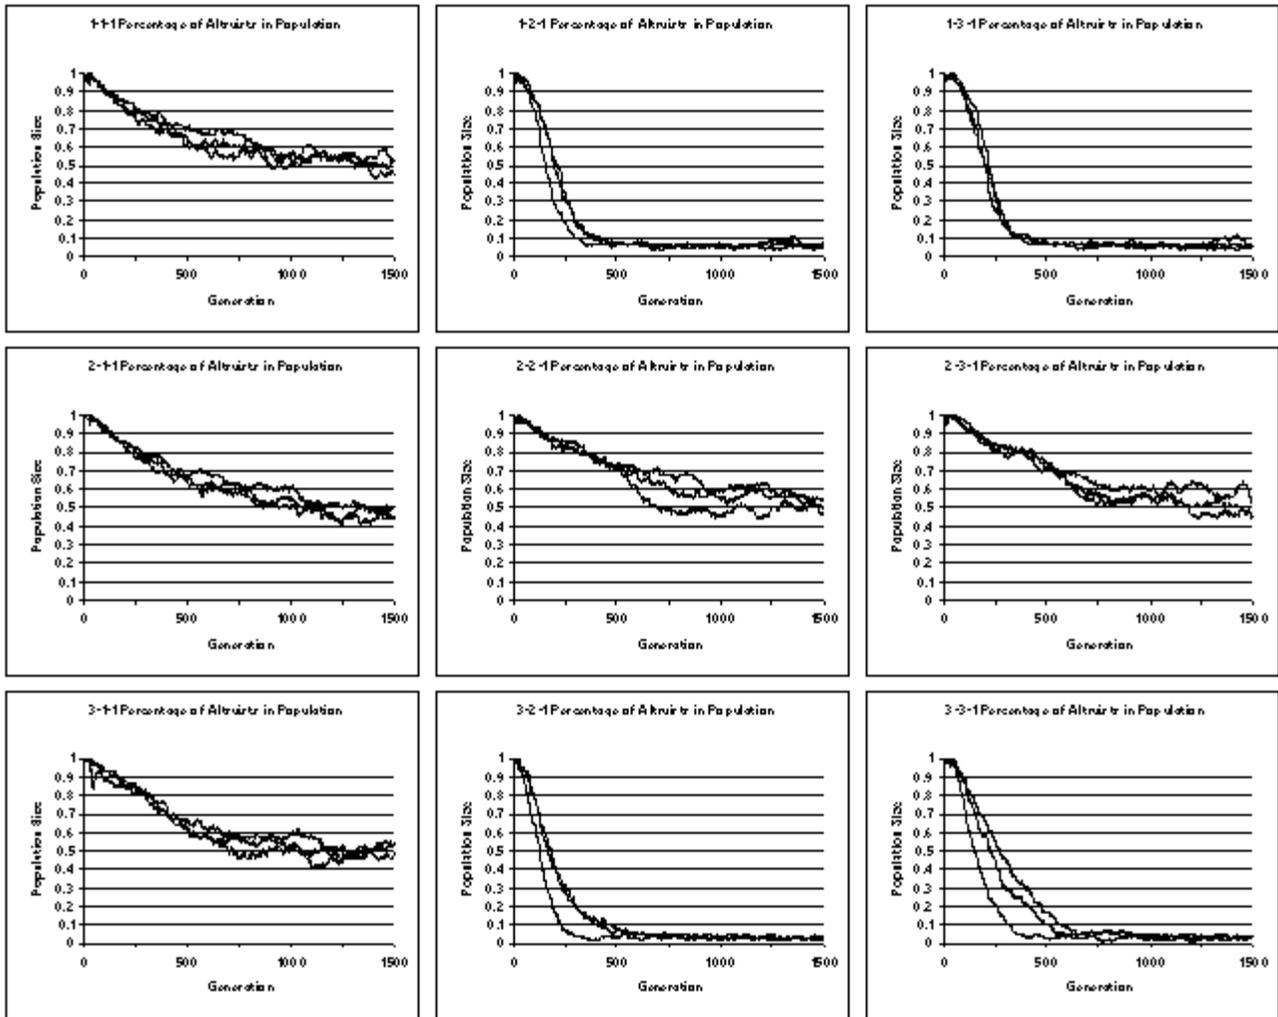


Figure 3: Percentage of altruists in the population with respect to initial levels (1=100%)

5 Discussion

Both goals of this research, as stated in Section 1, are successfully met. The proposed algorithm has been implemented, and altruism has, indeed, been shown to be selected and maintained by evolution in a number of cases. No direct comparison with the Barton's work could be made as his detailed results were not available in a suitable form. However, a few main points can be made. Firstly, it has been confirmed that the policy of progressive taxation produces more altruists than communism. An additional policy (poll tax) was studied in this research, which also introduced the new dimension of 'knowledge of the degree of kinship' in the experimental setup. Unlike Barton's, these experiments produced populations of

Table 2: Percentage of altruistic individuals in the population (1=100%). (Columns, from left to right: Royalty, Prediction and Unknown models of kinship recognition. Rows, top to bottom: Communism, Progressive Taxation and Poll Tax sharing functions.)



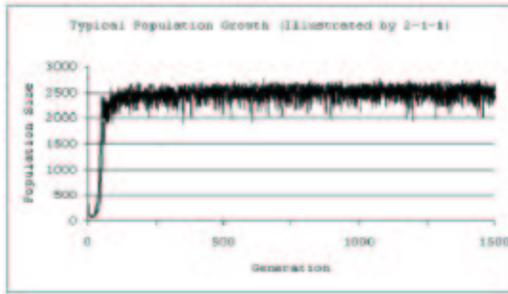


Figure 4: Evolution of population size

virtually the same size. Barton treats altruists and non-altruists as two different species, which in turn results in one species taking completely over the other one. In our results, there are several cases in which a balance between altruists and selfish individuals is maintained.

Altruism is a demonstration of the mechanisms on which natural selection is based. Note that this work does not aim to imply the existence of such gene in reality, and indeed nothing of the said above would change if one assumed altruistic behaviour being inherited not as a gene, but through upbringing.

There is interest in the use of natural selection in artificial societies. This research should bring the implementation of natural selection in artificial societies a step closer to the original mechanism that is copied. The authors' expectations are that the natural selection incorporating altruism would be suitable in cases, when the task is to produce an optimal population of agents rather than a single best individual, in situations when the knowledge about the performance of the population is incomplete and local.

The software described here may also represent a useful tool for the simulation of natural societies and give an interesting insight in their inner making, although this would be to experts in the relevant fields to judge.

The main characteristics of the model of altruism discussed here, namely, 'inherited' and 'kinship-driven', also mark the limits of its reach.

Firstly, the model does not allow changes in the altruistic behaviour of an individual within its lifespan. In fact, natural selection and individual learning are not perceived here as mutually exclusive. It is expected that, in many cases, combination of the two could be a successful strategy, where natural selection provides the starting point for the individual behaviour, which is modified according to the agent's personal experience. The actual technique employed at this second stage could be, for instance, based on game theory, where natural selection provides a suitable initial strategy. If individual behaviour is to be modified by a machine learning technique, natural selection could also provide it with a suitable bias. Research in this direction should be helped by the York MAS, currently under development, which supports natural selection among agents, as well as logic-based programming

of behaviour and individual learning (Kazakov and Kudenko, 2001).

The second limitation of the model of altruism discussed here is that it does not discuss the case when agents can at will opt in and out of a society promoting altruism among its members. Since the names of many such societies draw analogies with kinship, e.g. 'fraternities' or 'sororities', in order to evoke the corresponding spirit of altruism (or 'brotherhood') in its members, the authors believe that also in this case the findings described in the paper would not be without relevance.

In comparison with logic-based approaches, this research makes one simple initial assumption, and attempts to see if altruism can be worked out from first principles. The actual behaviour of agents can be deterministic (and described in logic) or stochastic, that should not be of principle importance. On the other hand, no further background knowledge is assumed here—the agent's rules of behaviour are let to evolve, and not set in advance. In the future, comparisons with Hamilton's analytical model, and the evolutionary game theory point of view would also be worth exploring.

6 Future Work

It would also be interesting to extend the platform developed to implement different mating policies, so that pairs of individuals could be selected from a single mating pool or from separate mating pools into which individuals have previously been grouped according to their internal state: rich meet (mostly) rich, poor meet poor, individuals with high sexual drive are grouped together, etc.

In addition, two environmental parameters, resource availability and the probability of meeting another individual, should be taken into account, and used to test the effectiveness of altruistic vs. selfish policy in various, and changing, environments.

An important and, potentially, non-trivial issue is the analysis of the content of the individuals' sets of genes and their evolution in time. In the case when the propagation of all genes is subject to simultaneous selection, one would have to study data sets, which are multidimensional—one dimension per locus plus an extra dimension representing time—hence difficult to visualise. One could expect that there would be a correlation between the genes selected in each locus, and that certain combinations might show a trend of dominating the population, which would form clusters around those points. Methods and tools for multivariate data visualisation with a minimal loss of information, such as those described by Schröder and Noy (2001), will be considered for the above task.

Acknowledgements

The second author wishes to express his gratitude to his wife María Elena and daughter Maia for being such a wonderful source of inspiration.

References

- John Barton. Kinship-driven altruism in multi-agent systems. Project report for a degree in Computer Science, University of York. Project supervisor: Dimitar Kazakov, 2001.
- David E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, 1989.
- W. D. Hamilton. The genetical evolution of social behaviour I. *Journal of Theoretical Biology*, 7:1–16, 1964a.
- W. D. Hamilton. The genetical evolution of social behaviour II. *Journal of Theoretical Biology*, 7:17–52, 1964b.
- D. Kazakov and D. Kudenko. *Multi-Agent Systems and Applications*, chapter Machine learning and inductive logic programming for multi-agent systems, pages 246–270. LNAI 2086. Springer, 2001.
- Michael Schröder and Penny Noy. Multi-agent visualisation based on multivariate data. In *Working Notes of the Fourth UK Workshop on Multi-Agent Systems UKMAS 2001*. 2001.
- Della Thompson, editor. *The Oxford Compact English Dictionary*. Oxford University Press, 1996.
- Tim Watson. Kin selection and cooperating agents. Technical report, Dept. of Computer Science, De Montfort University, Leicester, 1995.
- M. Wooldridge and Nick Jennings. Intelligent agents: theory and practice. *Knowledge Engineering Review*, 2(10), 1995.